

SYSTEM AND METHOD OF VIRTUAL PRIVATE NETWORK
ROUTE TARGET FILTERING

RELATED APPLICATION

The present application claims priority to provisional application Serial No. 60/294,755, filed on May 31, 2001, entitled "SYSTEM AND METHOD OF VIRTUAL PRIVATE NETWORK ROUTE TARGET FILTERING."

TECHNICAL FIELD OF THE INVENTION

This invention relates to telecommunications network and equipment, and more particularly, to a system and method of virtual private network route target filtering.

BACKGROUND OF THE INVENTION

Request for Comment (RFC) 2547bis provides out a virtual private network (VPN) model that uses border gateway protocol (BGP) to distribute VPN routing information across the service provider's backbone and Multi-protocol label switching (MPLS) to forward VPN traffic from one VPN site to another. RFC 2547bis defines a VPN as a collection of policies, and these policies control connectivity among a set of sites. A customer site is connected to the service provider network by one or more ports, where the service provider associates each port with a VPN routing table. RFC 2547bis calls the VPN routing table as a VPN routing and forwarding (VRF) table. A customer edge (CE) device provides customer access to the service provider network over a data link to one or more provider edge (PE) routers. The CE device can be a host, a Layer 2 switch, or more commonly, an IP router that establishes an adjacency with its directly connected PE routers. After the adjacency is established, the CE router advertises the site's local VPN routes to the PE router and learns remote VPN routes from the PE router. After learning local VPN routes from CE routers, a PE router exchanges VPN routing information with other PE routers using IBGP.

A route distinguisher (RD) is an identifier that is used to differentiate IP addresses or IPv4 prefixes of a VPN from another because customers may not use globally unique IP addresses. RFC 2547bis constrains the distribution of routing information among PE routers by the use of route filtering based on a route target (RT) attribute, which is one of the BGP extended community attributes. Route targets include import targets and export targets. The import target of a site governs which sites' route update information or advertisement it will accept; the export target of the site specifies what import target the sites it advertises to should include.

An enterprise's VPN may be configured in a hub-and-spoke topology where the firewall is the hub through which all traffic is routed. The hub site's VRF table is configured with an export target = hub and an import target = spoke. The VRF table at the hub site distributes all of the routes in its VRF table with a hub attribute that causes the routes to be imported by the spoke sites. The VRF table at the hub site imports all remote routes with a spoke attribute. The VRF table at each spoke site is configured with an export target = spoke and an import target = hub. The VRF table at each spoke site distributes its routes with a spoke attribute, which causes the routes to be imported by the hub site, but dropped by other spoke sites. The VRF table at a spoke site imports only routes with a hub attribute, which causes its VRF table to be populated only with routes advertised by the hub site.

In conventional VPNs, policy-based routing around the firewall in the hub-and-spoke topology requires either the knowledge of the IPv4 prefix or the use of at least two router ports in order to route packets to the spokes from the hub. The reliance using the IP address is tedious and labor intensive, and using an extra router port is inefficient and costly.

SUMMARY OF THE INVENTION

In accordance with an embodiment of the present invention, a system of route target filtering includes an import filter receiving a plurality of routes having a next hop routing information. The import filter accepts a first subset of the routes according to an import target policy. The system further includes a re-export filter also receiving the plurality of routes. The re-export filter modifies the next hop information of a second subset of the routes, and distributes the modified routes. If

desired, the re-export filter may also modify the RD and RT information of the second subset of the routes.

In accordance with another embodiment of the present invention, a network includes a hub node, and a plurality of spoke nodes in communications with one another via the hub node. The hub node includes an import filter receiving a plurality of routes. The plurality of routes each has a next hop routing information. The import filter accepts a first subset of the routes according to an import target policy. The network also includes a re-export filter receiving the plurality of routes. The re-export filter modifies the next hop information of a second subset of the routes, and distributes the modified routes. If desired, the re-export filter may also modify the RD and RT information of the second subset of the routes.

In accordance with yet another embodiment of the present invention, a method includes the steps of receiving a plurality of routes each having a next hop routing information, accepting a first subset of the plurality of routes according to a predetermined policy, modifying the next hop information of a second subset of the plurality of routes, and distributing the modified routes.

The present invention uses re-export filters to modify the advertised routes and sends the modified routes to the route reflector for distribution. The routes of nodes within the VPN is modified to have the next hop information as designating the firewall node. The present invention thus provides a way to perform policy routing around a firewall for a VPN based on RFC 2547bis without the disadvantages associated with the use of IP prefix knowledge or an extra port.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention, the objects and advantages thereof, reference is now made to the following descriptions taken in connection with the accompanying drawings in which:

FIGURE 1 is a simplified block diagram of a virtual private network (VPN) configured with a firewall according to the teachings of the present invention; and

FIGURE 2 is a simplified diagram of an embodiment of the route target filtering scheme according to the teachings of the present invention.

DETAILED DESCRIPTION OF THE DRAWINGS

The preferred embodiment of the present invention and its advantages are best understood by referring to FIGURES 1 and 2 of the drawings, like numerals being used for like and corresponding parts of the various drawings.

FIGURE 1 is a simplified block diagram of a virtual private network (VPN) 10 configured in a hub-and-spoke configuration with a firewall as the hub 12 according to the teachings of the present invention. Hub 12 is a customer edge (CE) device of Site₁ 16, which is coupled to a provider edge (PE) device 14 of a provider's network. A site is defined by Request for Comment (RFC) 2547bis as a collection of customer routers with IP connectivity. Virtual private network 10 is a set of sites, including Site₁ 16, Site₂ 18, Site₃ 20, and Site₄ 22. Provider edge device 14 may be directly or indirectly coupled to customer edge devices 24-28 of Site₂ 18, Site₃ 20, and Site₄ 22, respectively. Communication to and from an extranet 20 is done via customer edge device CE₁ 12, which is the firewall. According to the teachings of the present invention, this is done using a route target filtering scheme which does not have the disadvantages of conventional methods.

FIGURE 2 is a simplified diagram of an embodiment of the route target filtering scheme according to the teachings of the present invention. The hub, CE₁ 12, receives routes from a route reflector (not shown), which is a centralized distributor of routes. The route information includes the route distinguisher (RD), route target (RT), and next hop (NH):

$$RouteInformation = \{RD, RT, NH\}$$

Next hop information may include the provider edge external virtual private router IP address and the IfIndex. Other route information such as the site of origin, the VPN identifier and the IPv4 prefix may also be included in the route. A set of import filters 40 is used to determine which routes should be accepted and which routes should be rejected. Import filters 40 include a mask used to compare the route input to certain route information such as route target values. For example, the mask is used to indicate which route information field should be compared to the target value:

Mask{0|1,0|1,0|1}, Value{,*,*}, Action=accept|discard*

If the mask is set or one for a specific field, then the corresponding value for that field is compared with the received route; if the mask is clear or zero for a specific field, then the corresponding value for that field is not compared. Upon a match between the route information and the compare value, the route is either rejected or accepted. The accepted routes are passed on for route advertisement.

The filtering scheme also includes export filters, local export filters 42 and re-export filters 44. Local export filters 42 perform port level-based VPN assignments. Local export filters 42 receive routes from the PE-CE routing protocol and apply at least one filter. The accepted routes are exported to the proper route reflector. Re-export filters 44 also receive routes from the route reflector as input. The accepted routes are modified with a different route distinguisher, route target, and next hop information and redistributed to the route reflector.

As an illustrative example, customer edge device, CE₁ 12, has import targets RT_R, RT_S and RT_T. Its spokes, CE₂, CE₃ and CE₄, each respectively advertises export route targets RT_R, RT_S and RT_T. Sites 16-22 belong to a VPN that is distinguished by route distinguisher RD₁, for example. Therefore, route information advertised by CE₂ to the hub, CE₁, for example is {RD, RT, NH} = {RD₁, RT_R, CE₂}. Re-export filter, upon receiving this route, modifies the route to be {RD, RT, NH} = {RD₂, RT_x, CE₁}, for example. One or more sites in extranet 30 may import routes with RT_x as the route target. It is led to believe that to communicate with site 18 would require it to communicate with CE₁ 12 because CE₁ is designated as the next hop in the route information. A different route distinguisher, RD₂, is attached to the modified route in order to avoid duplication in the route reflector.

In this manner, routes to sites within a VPN are advertised with the firewall node as the next hop, so that all communications are routed via the firewall. The present invention does not require the manipulation of the IPv4 prefix or the use of an extra router port at the provider edge device to route data to sites within the VPN and outside the VPN. This saves the labor intensive and tedious management and manipulation of the IPv4 prefix and the costs associated with the extra router port. As IP addresses are constantly changing, independence therefrom also provides added

benefits. The present invention may be applicable to other situations where redirected routes are needed.

5 While the invention has been particularly shown and described by the foregoing detailed description, it will be understood by those skilled in the art that various changes, alterations, modifications, mutations and derivations in form and detail may be made without departing from the spirit and scope of the invention.

10044106.011102